

The Garden-Hose Game

A New Model of Computation, and Application to Position-Based Quantum Cryptography

Harry Buhrman^{*}, Serge Fehr, Christian Schaffner^{**}, and Florian Speelman^{*}

Centrum Wiskunde & Informatica (CWI), The Netherlands
University of Amsterdam, The Netherlands

Abstract. We study position-based cryptography in the quantum setting. We examine a class of protocols that only require the communication of a single qubit and $2n$ bits of classical information. To this end, we define a new model of communication complexity, the garden-hose model, which enables us to prove upper bounds on the number of EPR pairs needed to attack such schemes. This model furthermore opens up a way to link the security of quantum position-based cryptography to traditional complexity theory.

1 Introduction

Background: Position-based (Quantum) Cryptography

The goal of *position-based cryptography* is to use the geographical position of a party as its only “credential”. For example, one would like to send a message to a party at a geographical position pos with the guarantee that the party can decrypt the message only if he or she is physically present at pos . The general concept of position-based cryptography was introduced by Chandran, Goyal, Moriarty and Ostrovsky [CGMO09].

A central task in position-based cryptography is the problem of *position-verification*. We have a *prover* P at position pos , wishing to convince a set of *verifiers* V_0, \dots, V_k (at different points in geographical space) that P is indeed at that position pos . The prover can run an interactive protocol with the verifiers in order to convince them. The main technique for such a protocol is known as distance bounding [BC94]. In this technique, a verifier sends a random nonce to P and measures the time taken for P to reply back with this value. Assuming that the speed of communication is bounded by the speed of light, this technique gives an upper bound on the distance of P from the verifier.

The problem of secure position-verification has been studied before in the field of wireless security, and there have been several proposals for this task ([BC94,SSW03,VN04,Bus04,CH05,SP05,ZLFW06,CCS06]). However, [CGMO09] shows that there exists no protocol for secure position-verification that offers security in the presence of *multiple colluding* adversaries. In other words, the set of verifiers cannot distinguish between the case when they are interacting with an honest prover at pos and the case when they are interacting with multiple colluding dishonest provers, none of which is at position pos .

The impossibility result of [CGMO09] relies heavily on the fact that an adversary can locally store all information he receives *and* at the same time share this information with other colluding adversaries, located elsewhere. Due to the no-cloning theorem, such a strategy will not work in the quantum setting, which opens the door to secure protocols that use quantum information. The quantum model was first studied by Kent et al. under the name of “quantum tagging” [KMSB06,KMS11]. Several

^{*} Supported by a NWO VICI grant and the EU 7th framework grant QCS.

^{**} Supported by a NWO VENI grant.

schemes were developed [KMS11,Mal10a,CFG⁺10,Mal10b,LL11] and proven later to be insecure. Finally in [BCF⁺11] it was shown that in general no unconditionally secure quantum position-verification scheme is possible. Any scheme can be broken using a double exponential amount of EPR pairs in the size of the messages of the protocol. Later, Beigi and König improved in [BK11] the double exponential dependence to single exponential making use of port-based teleportation [IH08,IH09].

Due to the exponential overhead in EPR pairs, the general no-go theorem does not rule out the existence of quantum schemes that are secure for all practical purposes. Such schemes should have the property that the protocol, when followed honestly, is feasible, but cheating the protocol requires unrealistic amounts of resources, for example EPR pairs or time.

Analyzing the Beigi-König Scheme

To this end, Beigi and König [BK11] proposed a position-verification scheme using mutually unbiased bases. They showed that if the colluding parties are not allowed to send quantum, but only classical information to each other, then a linear amount of entanglement is necessary to break the scheme. They left open whether more entanglement was needed. As a first contribution, we close this gap and show that a linear number of EPR pairs is also *sufficient* to break the scheme.

An Interesting Class of Schemes

Furthermore, we consider a class of schemes that only involve a single qubit, and $2n$ classical bits. Such schemes were first considered by Kent et al. [KMS11]. We focus on the one-dimensional set-up. The schemes easily generalize to three-dimensional space. The prover wants to convince the two verifiers, V_0 and V_1 , that he is at position pos on the line in between them. V_0 sends a qubit $|\phi\rangle$ prepared in a random basis to P . In addition, V_0 sends a string $x \in \{0,1\}^n$ and V_1 a $y \in \{0,1\}^n$ to P . All messages are timed such that they arrive at the same time at P 's claimed position. After receiving $|\phi\rangle$, x and y , P computes a predetermined Boolean function $f(x,y)$.¹ He sends $|\phi\rangle$ to V_0 if $f(x,y) = 0$ and to V_1 otherwise. V_0 and V_1 check that they receive the correct qubit in time corresponding to pos and measure the received qubit in the basis corresponding to which it was prepared. In order to cheat the scheme, we imagine two provers P_0 and P_1 on either side of the claimed position pos , who try to simulate the correct behavior of an honest P at pos .

The attack described in [KMS11] and the general no-go theorems from [BCF⁺11,BK11] imply that there is a strategy for P_0 and P_1 such that they can accomplish the following. P_0 receives $|\phi\rangle$, x and P_1 receives y . They are allowed to simultaneously send a single message to each other such that upon receiving that message they both know $f(x,y)$ and if $f(x,y) = 0$ then P_0 still has $|\phi\rangle$, otherwise P_1 has it in his possession. This teleportation-based cheating strategy however requires an exponential amount of EPR pairs (in n). We show in this paper that the number of EPR pairs required for such a protocol can be upper-bounded by a complexity measure that is related to the non-uniform space complexity of computing f . This complexity can sometimes be much smaller. For example, it follows that if $f(x,y)$ can be computed in logspace, then there is a cheating strategy that only requires a polynomial amount of entanglement. Our proof is inspired by *permutation branching programs* introduced by Barrington [Bar89] and a general technique to make log-space computations reversible [LMT97].

The motivation for considering this particular protocol for position-verification is the hope that for “complicated enough” functions $f(x,y)$, the amount of entanglement needed to successfully break the security of the protocol grows (at least) linearly in the bit length n of the classical strings x,y .

¹ We assume for simplicity that computation does not take any time.

If this intuition is true, it is a very interesting property of the protocol that we obtain a favorable relation between quantum and classical difficulty of operations in the following sense: if we increase the length of the classical inputs x, y , we require more *classical* computing power of the honest prover, whereas more *quantum* resources (in form of entangled states) are required by the adversary to break the protocol. Thus, the more classical resources the honest users use to faithfully execute the scheme, the more quantum resources the adversary needs in order to break it. To the best of our knowledge, such a trade-off has never been observed for a quantum-cryptographic protocol.

We give some first indications that the above may indeed be true. We show that if f is injective in x (meaning that $\forall x \neq x' \exists y : f(x, y) \neq f(x', y)$) or in y (defined accordingly), then for any attack that succeeds with certainty, the two dishonest provers require a joint quantum working space consisting of at least a logarithmic amount of qubits in n . Also, we show that if the entangled starting state for the dishonest provers is *fixed*, e.g. a list of EPR pairs, then there exists a function f for which the starting state must consist of at least linearly many qubits in n to allow for a perfect attack. Restricting to perfect attacks makes the claims rather weak from a cryptographic point of view; we hope that this can be improved in future work.

The Garden-Hose Complexity

In order to isolate the properties of attacks on these one-qubit schemes, we define a new model of communication complexity which we call the *garden-hose model*. Alice and Bob as usual have to compute a Boolean function $f(x, y)$. In order to do so they possess a number of water pipes that lay between them. Moreover, they each have additional pieces of hose that they can use to connect up the ends of the water pipes that are at their side. For example, Alice may choose to connect pipe 17 with 19 and pipe 28 with 687 etc. Bob connects up the ends of the pipes on his side. For each input they can use a different connection scheme. In order to compute the function, Alice in addition has a source of water that she connects to one of the the pipes on her side. She now opens the water tap. It is easy to see that the water will flow out on one side only. If this is Alice's then they proclaim the function value to be 0 otherwise the function value is 1. We define the garden-hose complexity of f to be the minimum number of pipes needed to compute f .

The garden-hose model links the number of EPR pairs sufficient to attack a quantum position-verification scheme to traditional complexity theory: the number of EPR pairs needed for a successful attack is upper bounded by the garden-hose complexity of f . Unfortunately, so far it is unclear whether the garden-hose complexity by any means gives a lower bound on the number of EPR pairs needed. If it does, then this gives a nice handle on proving security of such schemes based on complexity-theoretical assumptions. In order to have a practical scheme, we will need a function f in the complexity class P that has “large” garden-hose complexity. The existence of a function in P with super-polynomial garden-hose complexity will separate the complexity class P from LOGSPACE, which is a long standing open problem.

Beyond its connection to position-based quantum cryptography, we feel that the garden-hose complexity is interesting in its own right, and trying to understand its connections to other complexity measures appears like a challenging goal. In this paper we give some first answers, but many questions regarding the garden-hose complexity require further research.

Summary

In summary, the main results of this paper are the following:

- We show that a quantum position-verification protocol by Beigi and König [BK11] can be attacked with a linear amount of EPR pairs, establishing that their lower bound is optimal up to a constant factor.

- We study an interesting class of position-verification schemes that may have the following property: the more classical resources the honest users use to faithfully execute the scheme, the more quantum resources the adversary needs in order to break it. We give some first results towards proving this desirable property.
- We introduce a new model of communication complexity, called the garden-hose model. The model is an abstraction of certain types of attacks against the above class of position-verification schemes. As such, tools from classical communication complexity can be used to obtain upper bounds on the number of EPR pairs needed to break a given scheme.
- We prove almost-linear lower bounds in the garden-hose model for concrete functions like inner product, majority, and equality. We show that random functions have exponential garden-hose complexity.
- We establish that all functions computable in log space have polynomial garden-hose complexity. As a corollary, we obtain the following interesting connection between proving the security of quantum protocols and classical complexity theory: If there is an f in P such that there is no attack on our scheme using a polynomial number of EPR pairs, then $P \neq \text{LOGSPACE}$.
- Our approach may lead to practical secure quantum position-verification schemes whose security is based on classical complexity-theoretical assumptions such as P is different from LOGSPACE .

2 Preliminaries

We assume that the reader is familiar with basic concepts of quantum information theory. We refer to [NC00] for an introduction and merely fix some notation here.

2.1 Quantum Teleportation

An important example of a 2-qubit state is the *EPR pair*, which is given by $|\Phi\rangle_{AB} = (|0\rangle_A|0\rangle_B + |1\rangle_A|1\rangle_B)/\sqrt{2} \in \mathcal{H}_A \otimes \mathcal{H}_B = \mathbb{C}^2 \otimes \mathbb{C}^2$ and has the following properties: if qubit A is measured in the computational basis, then a uniformly random bit $x \in \{0, 1\}$ is observed and qubit B collapses to $|x\rangle$. Similarly, if qubit A is measured in the Hadamard basis, then a uniformly random bit $x \in \{0, 1\}$ is observed and qubit B collapses to $H|x\rangle$.

The goal of quantum teleportation is to transfer a quantum state from one location to another by only communicating classical information. Teleportation requires pre-shared entanglement among the two locations. To teleport a qubit Q in an arbitrary unknown state $|\psi\rangle_Q$ from Alice to Bob, Alice performs a Bell-measurement on Q and her half of an EPR pair, yielding a classical measurement outcome $k \in \{0, 1, 2, 3\}$. Instantaneously, the other half of the corresponding EPR pair, which is held by Bob, turns into the state $\sigma_k|\psi\rangle$, where $\sigma_0, \sigma_1, \sigma_2, \sigma_3$ denote the four Pauli-corrections $\{\mathbb{I}, X, Z, XZ\}$, respectively. The classical information k is then communicated to Bob who can recover the state $|\psi\rangle$ by performing σ_k on his EPR half.

3 On the (In)Security of a Proposed Protocol For Position Verification

3.1 Mutually Unbiased Bases

We use the following standard definition of mutually unbiased bases.

Definition 3.1. *Two orthonormal bases $\{|e_i^a\rangle\}_{i=1,\dots,d}$ and $\{|e_j^b\rangle\}_{j=1,\dots,d}$ of \mathbb{C}^d are called mutually unbiased, if $|\langle e_i^a | e_j^b \rangle|^2 = \frac{1}{d}$ holds for all $i, j \in \{1, \dots, d\}$.*

A *Pauli operator* on an n -qubit state is the tensor product of n one-qubit Pauli matrices. Hence, there are 4^n Pauli operators in total. For $i \in \{0, 1, 2, 3\}^n$, we can write the Pauli operator O_i as

$$O_i = \sigma_{i_1}^1 \sigma_{i_2}^2 \dots \sigma_{i_n}^n = \prod_{k=1}^n \sigma_{i_k}^k$$

where σ_j^k is the j -th Pauli matrix acting on qubit k (tensored with the identity on the other qubits).

Excluding the identity, there are $4^n - 1$ Pauli operators. These can be partitioned in $2^n + 1$ distinct subsets consisting of $2^n - 1$ *commuting* operators each [LBZ02]. The 2^n common eigenvectors of such a set of $2^n - 1$ commuting operators define an orthonormal basis. It can be shown that for any such partitioning, the resulting $2^n + 1$ bases are pairwise mutually unbiased [LBZ02]. We denote by $|e_x^a\rangle$ the x -th basis vector of the a -th mutually unbiased basis of this construction, where $x \in \{0, 1\}^n$ and $a \in \mathcal{A}$ for a set \mathcal{A} of $2^n + 1$ elements.

In the following, we will exploit a special property of this construction of mutually unbiased bases in order to attack a protocol for position-verification recently proposed by Beigi and König [BK11]. In particular, we use the fact that applying a Pauli operator only permutes the basis vectors *within* every mutually unbiased basis, but does not map any basis vector into another basis. This property is captured by the following lemma.

Lemma 3.2. *Let U be an arbitrary Pauli operator on n qubits. For arbitrary $a \in \mathcal{A}$ and $x \in \{0, 1\}^n$, let $|e_x^a\rangle$ be the x -th basis vector of the a -th mutually unbiased basis obtained from the construction above. Then, there exists $z \in \{0, 1\}^n$ such that $U|e_x^a\rangle = |e_z^a\rangle$.*

Proof. We can write U as

$$U = \sigma_{r_1}^1 \sigma_{r_2}^2 \dots \sigma_{r_n}^n = \prod_{k=1}^n \sigma_{r_k}^k.$$

Assume $|e_x^a\rangle$ is a common eigenvector of an internally commuting subset A of the Pauli operators, like described earlier. Denote the $2^n - 1$ elements of A by O_ℓ^A with $\ell \in \{1, \dots, 2^n - 1\}$. Note that $\sigma_0 \sigma_i = \sigma_i \sigma_0$ for $i \in \{0, 1, 2, 3\}$ and $\sigma_i \sigma_j = (-1)^{\delta_{ij}} \sigma_j \sigma_i$ for $i, j \in \{1, 2, 3\}$ and δ_{ij} the Kronecker δ -function. Because $|e_x^a\rangle$ is a common eigenvector of the Pauli operators in this set, it holds for all ℓ that $O_\ell^A |e_x^a\rangle = \lambda_\ell |e_x^a\rangle$ for some eigenvalue λ_ℓ . To prove the claim, we show that $U|e_x^a\rangle$ is also an eigenvector of all O_ℓ^A , with some (possibly different) eigenvalue λ'_ℓ .

$$\begin{aligned} O_\ell^A U |e_x^a\rangle &= \prod_{k=1}^n \sigma_{\ell_k}^k \sigma_{r_k}^k |e_x^a\rangle \\ &= (-1)^{\alpha(r, \ell)} \prod_{k=1}^n \sigma_{r_k}^k \sigma_{\ell_k}^k |e_x^a\rangle \\ &= (-1)^{\alpha(r, \ell)} U O_\ell^A |e_x^a\rangle \\ &= \lambda'_\ell U |e_x^a\rangle, \end{aligned}$$

where we define $\lambda'_\ell := (-1)^{\alpha(r, \ell)} \lambda_\ell$ and the function $\alpha(r, \ell)$ determines the phase arising from the commutation relations of the σ_{r_k} 's and σ_{ℓ_k} 's. Because $U|e_x^a\rangle$ is a common eigenvector of all O_ℓ^A , there exists $z \in \{0, 1\}^n$ such that $|e_z^a\rangle = U|e_x^a\rangle$. \square

3.2 The Protocol

The protocol described in Figure 1 uses an (almost) complete set of mutually unbiased bases $\{|e_x^a\rangle_{x=1,\dots,2^n}\}^{a \in \{0,1\}^n}$ as defined above. The protocol can be seen as a higher-dimensional extension of the basic BB84-protocols proposed and analyzed in [KMS11, BCF⁺11]. In [BK11], Beigi and König show that PV_{MUB} is secure against adversaries that share fewer than $n/2$ EPR pairs and are restricted to one round of simultaneous *classical* communication. They leave open whether the protocol remains secure against colluding adversaries that share more entanglement. We answer this question here. In the rest of the section, we show that for the construction of MUBs mentioned above, it is sufficient for adversaries to share n EPR pairs in order to perfectly break the protocol PV_{MUB} . It follows that the lower bound on the number of EPR pairs given in [BK11] is optimal up to constant factors.

0. V_0 and V_1 share common (secret) randomness in the form of uniformly distributed bitstrings $a, x \in \{0,1\}^n$.
1. V_0 sends a to P and V_1 prepares the state $|e_x^a\rangle$ and sends it to P . The timing is chosen such that both the classical information and the quantum state arrive at the prover at the same time.
2. P measures the state in the basis $\{|e_i^a\rangle\}_i$, getting measurement outcome $\hat{x} \in \{0,1\}^n$. He sends \hat{x} to both V_0 and V_1 .
3. V_0 and V_1 accept if they receive \hat{x} at times consistent with \hat{x} being emitted from the claimed position in both directions simultaneously, and $\hat{x} = x$.

Fig. 1. Protocol PV_{MUB} from [BK11] for position-verification using mutually unbiased bases.

3.3 The Attack

The attack reported here is very similar to the attack on the BB84-scheme described in [KMS11]. The colluding adversaries \tilde{P}_0 en \tilde{P}_1 set up between the prover's claimed position and the verifiers V_0 and V_1 , intercepting messages from V_0 and V_1 .

Adversary \tilde{P}_0 has knowledge of the basis a and \tilde{P}_1 gets the state $|e_x^a\rangle$. Our attack shows that using n ebits and one round of simultaneous classical communication suffices to determine x , and thus breaking protocol PV_{MUB} . We assume that the set of mutually unbiased bases used is equivalent to a basis obtained by a partitioning of Pauli operators as described above. To the best of our knowledge, any currently known construction of mutually unbiased basis sets of dimension 2^n is of this form. If the used set of mutually unbiased bases differs from one of these by a unitary transform, the attack still works by the adversaries just applying this unitary before the first step.

As soon as \tilde{P}_1 receives the state $|e_x^a\rangle$, she teleports it to \tilde{P}_0 and forwards the classical outcome of the teleportation measurement indicating the needed Pauli correction U . Using Lemma 3.2, the teleported state is still a basis vector of the same mutually unbiased basis, i.e. the state \tilde{P}_0 has before correction is $|e_z^a\rangle$, with z depending on the teleportation measurement outcome. \tilde{P}_0 measures $|e_z^a\rangle$ in basis a , getting outcome z which she sends to \tilde{P}_1 .

Now both adversaries hold a , z and the teleportation Pauli correction U . It is straightforward to (classically) derive the correct x , it is the measurement outcome in basis a of the state $|e_x^a\rangle = U|e_z^a\rangle$ after applying the Pauli correction U .

4 The Garden-Hose Game

4.1 Motivation

The results of this section are motivated by the study of a particular quantum protocol for secure position verification, described in Figure 2. The protocol is of the generic form described in Section 3.2 of [BCF⁺11]. In Step 0, the verifiers prepare challenges for the prover. In Step 1, they send the challenges, timed in such a way that they all arrive at the same time at the prover. In Step 2, the prover computes his answers and sends them back to the verifiers. Finally, in Step 3, the verifiers verify the timing and correctness of the answer.

As in [BCF⁺11], we consider here for simplicity the case where all players live in one dimension, the basic ideas generalize to higher dimensions. In one dimension, we can focus on the case of two verifiers V_0, V_1 and an honest prover P in between them.

We minimize the amount of quantum communication in that only one verifier, say V_0 , sends a qubit to the prover, whereas both verifiers send classical n -bit strings $x, y \in \{0, 1\}^n$ that arrive at the same time at the prover. We fix a publicly known boolean function $f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$ whose output $f(x, y)$ decides whether the prover has to return the qubit (unchanged) to verifier V_0 (in case $f(x, y) = 0$) or to verifier V_1 (if $f(x, y) = 1$).

0. V_0 randomly chooses two n -bit strings $x, y \in \{0, 1\}^n$ and privately sends y to V_1 . V_0 prepares an EPR pair $(|0\rangle_V |0\rangle_P + |1\rangle_V |1\rangle_P) / \sqrt{2}$. If $f(x, y) = 0$, V_0 keeps the qubit in register V . Otherwise, V_0 sends the qubit in register V privately to V_1 .
1. V_0 sends the qubit in register P to the prover P together with the classical n -bit string x . V_1 sends y so that it arrives at the same time as the information from V_0 at P .
2. P evaluates $f(x, y) \in \{0, 1\}$ and routes the qubit to $V_{f(x, y)}$.
3. V_0 and V_1 accept if the qubit arrives in time at the right verifier and the Bell measurement of the received qubit together with the qubit in V yields the correct outcome.

Fig. 2. Position-verification scheme PV_{qubit} using one qubit and classical n -bit strings.

The motivation for considering this protocol is the following: As the protocol uses only one qubit which needs to be correctly routed, the honest prover’s quantum actions are trivial to perform. His main task is evaluating a classical boolean function f on classical inputs x and y whose bit size n can be easily scaled up. On the other hand, our results in this section suggest that the adversary’s job of successfully attacking the protocol becomes harder and harder for larger input strings x, y . The hope is that for “complicated enough” functions $f(x, y)$, the amount of EPR pairs (ebits) needed to successfully break the security of the protocol PV_{qubit} grows (at least) linearly in the bit length n of the classical strings x, y .

If this intuition can be proven to be true, it is a very interesting property of the protocol that we obtain a favorable relation between quantum and classical difficulty of operations in the following sense: if we increase the length of the classical inputs x, y , we require more *classical* computing power of the honest prover, whereas more *quantum* resources (ebits) are required by the adversary to break the protocol. To the best of our knowledge, such a trade-off has never been observed for a quantum-cryptographic protocol.

In order to analyze the security of the protocol PV_{qubit} , we define the following communication game in which Alice and Bob play the roles of the adversarial attackers of PV_{qubit} . Alice starts with an unknown qubit $|\phi\rangle$ and a classical n -bit string x while Bob holds the n -bit string y . They also share

some quantum state $|\eta\rangle_{AB}$ and both players know the Boolean function $f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$. The players are allowed one round of simultaneous classical communication combined with arbitrary local quantum operations. When $f(x, y) = 0$, Alice should be in possession of the state $|\phi\rangle$ at the end of the protocol and on $f(x, y) = 1$, Bob should hold it.

As a simple example consider the case where $f(x, y) = x \oplus y$, the exclusive OR function, with 1-bit inputs x and y . Alice and Bob then have the following way of performing this task perfectly by using a pre-shared quantum state consisting of three EPR pairs (three ebits). Label the first two EPR pairs 0 and 1. Alice teleports $|\phi\rangle$ to Bob using the pair labeled with her input x . This yields measurement result $i \in \{0, 1, 2, 3\}$, while Bob teleports his half of the EPR pair labeled y to Alice using his half of the third EPR pair while obtaining measurement outcome $j \in \{0, 1, 2, 3\}$. In the round of simultaneous communication, both players send the classical measurement results and their inputs x or y to the other player. If $x \oplus y = 1$, i.e. x and y are different bits, Bob can apply the Pauli operator σ_i to his half of the EPR pair labeled $x = y \oplus 1$, correctly recovering $|\phi\rangle$. Similarly, if $x \oplus y = 0$, it is easy to check that Alice can recover the qubit by applying $\sigma_i \sigma_j$ to her half of the third EPR pair.

If Alice and Bob are *constrained* to the types of actions in the example above, i.e., if they are restricted to teleporting the quantum state back and forth depending on their classical inputs, we obtain the following notion of garden-hose game and garden-hose complexity.

4.2 Definition of the Garden-Hose Game

Alice and Bob get n -bit input strings x and y , respectively. Their goal is to “compute” an agreed-upon Boolean function $f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$ on these inputs, in the following way. We assume that Alice and Bob have s pipes between them. Depending on their respective classical inputs x and y , they connect their ends of the pipes with pieces of hose, of which they have an unlimited amount. Note however, that we do not allow “T-pieces” (or more complicated constructions) of hose which connect two or more pipes to one, or vice versa; only one-to-one connections are allowed. Alice has a source of water which she connects to one of the pipes, and then she turns on the water. It is easy to check that no “deadlocks” are possible and hence the water will flow out on either of the sides. They succeed in computing f (we may also say: they win the garden-hose game), if the water comes out of one of the pipes on Alice’s side whenever $f(x, y) = 0$, and the water comes out of one of the pipes on Bob’s side whenever $f(x, y) = 1$. Note that it does not matter out of which pipe the water flows, only on which side it flows. We stress once more that what makes the game non-trivial is that Alice and Bob must do their “plumbing” based on their local input only, and they are not allowed to communicate. We refer to Figure 3 for an illustration of computing the XOR function in the garden-hose model.

We can translate any strategy of Alice and Bob in the garden-hose game to a perfect quantum attack of PV_{qubit} by using one EPR pair per pipe and performing Bell measurements where the players connect the pipes. Our hope is that also the converse is true in spirit: if many pipes are required to compute f , say we need superpolynomially many, then the number of EPR pairs needed for Alice and Bob to successfully break PV_{qubit} with probability close to 1 by means of an *arbitrary* attack (not restricted to Bell measurements on EPR pairs) should also be superpolynomial. We leave this as an interesting problem for future research. We stress that for this application, a polynomial lower bound on the number of pipes, and thus on the number of EPR pairs, is already interesting.

We formalize the above description of the garden-hose game, given in terms of pipes and hoses etc., by means of rigorous graph-theoretic terminology. However, we feel that the above terminology captures the notion of a garden-hose game very well, and thus we sometimes use the above “watery” terminology. We start with a balanced bi-partite graph $(A \cup B, E)$ which is 1-regular and where the cardinality of A and B is $|A| = |B| = s$, for an arbitrary large $s \in \mathbb{N}$. We slightly abuse notation and




Fig. 3. Garden-hose game for the XOR function.

denote both the vertices in A and in B by the integers $1, \dots, s$. If we need to distinguish $i \in A$ from $i \in B$, we use the notation i^A and i^B . We may assume that E consists of the edges that connect $i \in A$ with $i \in B$ for every $i \in \{1, \dots, s\}$, i.e., $E = \{\{i^A, i^B\} : 1 \leq i \leq s\}$. These edges in E are the *pipes* in the above terminology. We now extend the graph to $(A_\circ \cup B, E)$ by adding a vertex 0 to A , resulting in $A_\circ = A \cup \{0\}$. This vertex corresponds to the *water tap*, which Alice can connect to one of the pipes. Given a Boolean function $f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$, consider two functions E_{A_\circ} and E_B ; both take as input a string in $\{0, 1\}^n$ and output a set of edges (without self loops). For any $x, y \in \{0, 1\}^n$, $E_{A_\circ}(x)$ is a set of edges on the vertices A_\circ and $E_B(y)$ is a set of edges on the vertices B , so that the resulting graphs $(A_\circ, E_{A_\circ}(x))$ and $(B, E_B(y))$ have maximum degree at most 1. $E_{A_\circ}(x)$ consists of the *connections* among the pipes (and the tap) on Alice's side (on input x), and correspondingly for $E_B(y)$. For any $x, y \in \{0, 1\}^n$, we define the graph $G(x, y) = (A_\circ \cup B, E \cup E_{A_\circ}(x) \cup E_B(y))$ by adding the edges $E_{A_\circ}(x)$ and $E_B(y)$ to E . $G(x, y)$ consists of the pipes with the connections added by Alice and Bob. Note that the vertex $0 \in A_\circ$ has degree at most 1, and the graph $G(x, y)$ has maximum degree at most two 2; it follows that the maximal path $\pi(x, y)$ that starts at the vertex $0 \in A_\circ$ is uniquely determined. $\pi(x, y)$ represents the flow of the water, and the endpoint of $\pi(x, y)$ determines whether the water comes out on Alice or on Bob's side (depending on whether it is in A_\circ or in B).

Definition 4.1. A **garden-hose game** is given by a graph function $G : (x, y) \mapsto G(x, y)$ as described above. The number of pipes s is called the **size** of G , and is denoted as $s(G)$. A garden-hose game G is said to **compute** a Boolean function $f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$ if the endpoint of the maximal path $\pi(x, y)$ starting at 0 is in A_\circ whenever $f(x, y) = 0$ and in B whenever $f(x, y) = 1$.

Definition 4.2. The (deterministic) **garden-hose complexity** of a Boolean function $f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$ is the size $s(G)$ of the smallest garden-hose game G that computes f . We denote it by $GH(f)$.

We start with a simple upper bound on $GH(f)$ which is implicitly proven in the attack on Scheme II in [KMS11].

Proposition 4.3. For every Boolean function $f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$, the garden-hose complexity is at most $GH(f) \leq 2^n + 1$.

Proof. We identify $\{0, 1\}^n$ with $\{1, \dots, 2^n\}$ in the natural way. For $s = 2^n + 1$ and the resulting bipartite graph $(A_\circ \cup B, E)$, we can define E_{A_\circ} and E_B as follows. $E_{A_\circ}(x)$ is set to $\{(0, x)\}$, meaning that Alice connects the tap with the pipe labeled by her input x . To define E_B , group the set $Z(y) = \{a \in \{0, 1\}^n : f(a, y) = 0\}$ arbitrarily into disjoint pairs $\{a_1, a_2\} \cup \{a_3, a_4\} \cup \dots \cup \{a_{\ell-1}, a_\ell\}$ and set $E_B(y) = \{\{a_1, a_2\}, \{a_3, a_4\}, \dots, \{a_{\ell-1}, a_\ell\}\}$. If $\ell = |Z(y)|$ is odd so that the decomposition into pairs results in a left-over $\{a_\ell\}$, then a_ℓ is connected with the “reserve” pipe labeled by $2^n + 1$.

By construction, if $x \in Z(y)$ then $x = a_i$ for some i , and thus pipe $x = a_i$ is connected on Bob’s side with pipe a_{i-1} or a_{i+1} , depending on the parity of i , or with the “reserve” pipe, and thus $\pi(x, y)$ is of the form $\pi(x, y) = (0, x^A, x^B, v^B, v^A)$, ending in A_\circ . On the other hand, if $x \notin Z(y)$, then pipe x is not connected on Bob’s side, and thus $\pi(x, y) = (0, x^A, x^B)$, ending in B . This proves the claim. \square

We notice that the same proof shows that the garden-hose complexity $GH(f)$ is at most $2^k + 1$, when k is the one-way communication complexity from Alice to Bob of f .²

We introduce the following terminology. We say that a function $f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$ is *obtained* from a function $g : \{0, 1\}^m \times \{0, 1\}^m \rightarrow \{0, 1\}$ by *local pre-processing* if f is of the form $f(x, y) = g(\alpha(x), \beta(y))$, where α and β are arbitrary functions $\{0, 1\}^n \rightarrow \{0, 1\}^m$. The following invariance under local preprocessing follows immediately from the definition of the garden-hose complexity.

Lemma 4.4. *If f is obtained from g by local pre-processing, then $GH(f) \leq GH(g)$.*

4.3 Garden-Hose Complexity and Log-Space Computations

The following theorem shows that for a large class of functions, a polynomial amount of pipes suffices to win the garden-hose game. A function f with an n -bit input is log-space computable if there is a deterministic Turing machine M and a constant c , such that M outputs the correct value of f , and at most $c \cdot \log n$ locations³ of M ’s work tapes are ever visited by M ’s head during computation of every input of length n .

Theorem 4.5. *If $f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$ is log-space computable, then $GH(f)$ is polynomial in n .*

In combination with Lemma 4.4, it follows immediately that the same conclusion also holds for functions that are *log-space computable up to local pre-processing*, i.e., for any function $f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$ that is obtained from a log-space computable function $g : \{0, 1\}^m \times \{0, 1\}^m \rightarrow \{0, 1\}$ by local pre-processing, where m is polynomial in n . Below, in Proposition 4.7, we show that log-space up to local pre-processing is also *necessary* for a polynomial garden-hose complexity.

We will later see (Proposition 4.11) that there exist functions with large garden-hose complexity. However, a negative implication of Theorem 4.5 is that proving the existence of a *polynomial-time computable* function f with exponential garden-hose complexity is at least as hard as separating L from P , a long-standing open problem in complexity theory.

Corollary 4.6. *If there exists a function $f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$ in P that has superpolynomial garden-hose complexity, then $P \neq L$.*

Proof (of Theorem 4.5). Let M be a deterministic Turing machine deciding $f(x, y) = 0$. We assume that M ’s read-only input tape is of length $2n$ and contains x on positions 1 to n and y on positions $n + 1$ to $2n$. By assumption M uses logarithmic space on its work tapes.

² Or if needed, with a small adjustment in the protocol, $2^k + 2$ with k the one-way communication complexity of Bob to Alice.

³ All logarithms in this paper are with respect to base 2.

In this proof, a *configuration* of M is the location of its tape heads, the state of the Turing machine and the content of its work tapes, excluding the content of the read-only input tape. This is a slightly different definition than usual, where the content of the input tape is also part of a configuration. When using the normal definition (which includes the content of all tapes), we will use the term *total configuration*. Any configuration of M can be described using a logarithmic number of bits, because M uses logarithmic space.

A Turing machine is called *deterministic*, if every total configuration has a unique next one. A Turing machine is called *reversible* if in addition to being deterministic, every total configuration also has a unique predecessor. An $S(n)$ space-bounded deterministic Turing machine can be simulated by a reversible Turing machine in space $O(S(n))$ [LMT97]. This means that without loss of generality, we can assume M to be a reversible Turing machine, which is crucial for our construction. Let M also be *oblivious*⁴ in the tape head movement on the input tape. This can be done with only a small increase in space by adding a counter.

Alice's and Bob's perfect strategies in the garden-hose game are as follows. They list all configurations where the head of the input tape is on position n coming from position $n + 1$. Let us call the set of these configurations C_A . Let C_B be the analogous set of configurations where the input tape head is on position $n + 1$ after having been on position n the previous step. Because M is oblivious on its input tape, these sets depend only on the function f , but not on the input pair (x, y) . The number of elements of C_A and C_B is at most polynomial, being exponential in the description length of the configurations. Now, for every element in C_A and C_B , the players label a pipe with this configuration. Also label $|C_A|$ pipes **ACCEPT** and $|C_B|$ of them **REJECT**. These steps determine the number of pipes needed, Alice and Bob can do this labeling beforehand.

For every configuration in C_A , with corresponding pipe p , Alice runs the Turing machine starting from that configuration until it either accepts, rejects, or until the input tape head reaches position $n + 1$. If the Turing machine accepts, Alice connects p to the first free pipe labeled **ACCEPT**. On a reject, she leaves p unconnected. If the tape head of the input tape reaches position $n + 1$, she connects p to the pipe from C_B corresponding to the configuration of the Turing machine when that happens. By her knowledge of x , Alice knows the content of the input tape on positions 1 to n , but not the other half. Alice also runs M from the starting configuration, connecting the water source to a target pipe with a configuration from C_B depending on the reached configuration.

Bob connects the pipes labeled by C_B in an analogous way: He runs the Turing machine starting with the configuration with which the pipe is labeled until it halts or the position of the input tape head reaches n . On accepting, the pipe is left unconnected and if the Turing machine rejects, the pipe is connected to one of the pipes labeled **REJECT**. Otherwise, the pipe is connected to the one labeled with the configuration in C_A , the configuration the Turing machine is in when the head on the input tape reached position n .

In the garden-hose game, only one-to-one connections of pipes are allowed. Therefore, to check that the described strategy is a valid one, the simulations of two different configurations from C_A should never reach the same configuration in C_B . This is guaranteed by the reversibility of M as follows. Consider Alice simulating M starting from different configurations $c \in C_A$ and $c' \in C_A$. We have to check that their simulation can not end at the same $d \in C_B$, because Alice can not connect both pipes labeled c and c' to the same d . Because M is reversible, we can in principle also simulate M backwards in time starting from a certain configuration. In particular, Alice can simulate M backwards starting with configuration d , until the input tape head position reaches $n + 1$. The configuration of M at that

⁴ A Turing machine is called *oblivious*, if the movement in time of the heads only depend on the length of the input, known in advance to be $2n$, but not on the input itself. For our construction we only require the input tape head to have this property.

time can not simultaneously be c and c' , so there will never be two different pipes trying to connect to the pipe labeled d .

It remains to show that, after the players link up their pipes as described, the water comes out on Alice's side if M rejects on input (x, y) , and that otherwise the water exits at Bob's. We can verify the correctness of the described strategy by comparing the flow of the water directly to the execution of M . Every pipe the water flows through corresponds to a configuration of M when it runs starting from the initial state. So the side on which the water finally exits also corresponds to whether M accepts or rejects. \square

Proposition 4.7. *Let $f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$ be a Boolean function. If $GH(f)$ is polynomial (in n), then f is log-space computable up to local pre-processing.*

Proof. Let G be the garden-hose game that achieves $s(G) = GH(f)$. We write s for $s(G)$, the number of pipes, and we let E_{A_0} and E_B be the underlying edge-picking functions, which on input x and y , respectively, output the connections that Alice and Bob apply to the pipes. Note that by assumption, s is polynomial. Furthermore, by the restrictions on E_{A_0} and E_B , on any input, they consist of at most $(s + 1)/2$ connections.

We need to show that f is of the form $f(x, y) = g(\alpha(x), \beta(y))$, where α and β are arbitrary functions $\{0, 1\}^n \rightarrow \{0, 1\}^m$, $f : \{0, 1\}^m \times \{0, 1\}^m \rightarrow \{0, 1\}$ is log-space computable, and m is polynomial in n . We define α and β as follows. For any $x, y \in \{0, 1\}^n$, $\alpha(x)$ is simply a natural encoding of $E_{A_0}(x)$ into $\{0, 1\}^m$, and $\beta(y)$ is a natural encoding of $E_B(y)$ into $\{0, 1\}^m$. In the hose-terminology we say that $\alpha(x)$ is a binary encoding of the connections of Alice, and $\beta(y)$ is an encoding of the connections of Bob. Obviously, this can be done with m of polynomial size. Given these encodings, finding the endpoint of the maximum path $\pi(x, y)$ starting in 0 can be done with logarithmic space: at any point during the computation, the Turing machine only needs to maintain a couple of pointers to the inputs and a constant number of binary flags. Thus, the function g that computes $g(\alpha(x), \beta(y)) = f(x, y)$ is log-space computable in m and thus also in n . \square

4.4 Lower Bounds

In this section, we present lower bounds on the number of pipes required to win the garden-hose game for particular (classes of) functions.

Definition 4.8. *We call a function f injective for Alice, if for every two different inputs x and x' there exists y such that $f(x, y) \neq f(x', y)$. We define injective for Bob in an analogous way: for every $y \neq y'$, there exists x such that $f(x, y) \neq f(x, y')$ holds.*

Proposition 4.9. *If f is injective for Bob or f is injective for Alice, then*

$$GH(f) \log(GH(f)) \geq n.$$

Proof. We give the proof when f is injective for Bob. The proof for the case where f is injective for Alice is the same. Consider a garden-hose game G that computes f . Let s be its size $s(G)$. Since, on Bob's side, every pipe is connected to at most one other pipe, there are at most $s^s = 2^{s \log(s)}$ possible choices for $E_B(y)$, i.e., the set of connections on Bob's side. Thus, if $2^{s \log(s)} < 2^n$, it follows from the pigeonhole principle that there must exist y and y' in $\{0, 1\}^n$ for which $E_B(y) = E_B(y')$, and thus for which $G(x, y) = G(x, y')$ for all $x \in \{0, 1\}^n$. But this cannot be since G computes f and $f(x, y) \neq f(x, y')$ for some x due to the injectivity for Bob. Thus, $2^{s \log(s)} \geq 2^n$ which implies the claim. \square

We can use this result to obtain an almost linear lower bound for several functions that are often studied in communication complexity settings such as:

- Bit wise inner product: $\text{IP}(x, y) = \sum_i x_i y_i \pmod{2}$
- Equality: $\text{EQ}(x, y) = 1$ iff $x = y$
- Majority function: $\text{MAJ}(x, y) = 1$ iff $\sum_i x_i y_i \geq \lceil \frac{n}{2} \rceil$

For the first two functions we have a method that is linear in the number of pipes, and our best upper bound of computing majority in the garden-hose model is quadratic. Specific constructions can be found in [Spe11]. All of these functions are injective for both Alice and Bob, giving us the following corollary.

Corollary 4.10. *The functions bitwise inner product, equality and majority have garden-hose complexity at least $\frac{n}{\log(n)}$.*

Proposition 4.11. *There exist functions $f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$ for which $\text{GH}(f)$ is exponential.*

Proof. The existence of functions with an exponential garden-hose complexity can be shown by a simple counting argument. There are 2^{2^n} different functions $f(x, y)$. For a given size $s = s(G)$ of G , for every $x \in \{0, 1\}^n$, there are at most $(s + 1)^{s+1}$ ways to choose the connections $E_{A_o}(x)$ on Alice's side, and thus there are at most $((s + 1)^{s+1})^{2^n} = 2^{2^n(s+1)\log(s+1)}$ ways to choose the function E_{A_o} . Similarly for E_B , there are at most $2^{2^n s \log(s)}$ ways to choose E_B . Thus, there are at most $2^{2 \cdot 2^n(s+1)\log(s+1)}$ ways to choose G of size s . Clearly, in order for every function f to have a G of size s that computes it, we need that $2 \cdot 2^n(s+1)\log(s+1) \geq 2^{2^n}$, and thus that $(s+1)\log(s+1) \geq 2^{n-1}$, which means that s must be exponential. \square

5 Lower Bound In The Real World

In this section, we show that for a function that is injective for Alice or injective for Bob (according to Definition 4.8), the dimension of the state the adversaries need to handle (including possible quantum communication between them) in order to attack protocol PV_{qubit} perfectly has to be of order at least linear in the classical input size n . In other words, they require at least a logarithmic number of qubits in order to successfully attack PV_{qubit} . We start by showing two lemmas. The actual bound is shown in Section 5.3. In the last section, we show that there exist functions for which perfect attacks on PV_{qubit} requires the adversaries to handle a polynomial amount of qubits.

5.1 Localized Qubits

Assume we have two bipartite states $|\psi^0\rangle$ and $|\psi^1\rangle$ with the property that $|\psi^0\rangle$ allows Alice to locally extract a qubit and $|\psi^1\rangle$ allows Bob to locally extract the same qubit. Intuitively, these two states have to be different.

More formally, we assume that both states consist of five registers $R, A, \tilde{A}, B, \tilde{B}$ where registers R, A, B are one-qubit registers and \tilde{A} and \tilde{B} are arbitrary. We assume that there exist local unitary transformations $U_{A\tilde{A}}$ acting on registers $A\tilde{A}$ and $V_{B\tilde{B}}$ acting on $B\tilde{B}$ such that⁵

$$U_{A\tilde{A}}|\psi^0\rangle_{RA\tilde{A}B\tilde{B}} = |\beta\rangle_{RA} \otimes |P\rangle_{\tilde{A}B\tilde{B}} \quad (1)$$

$$V_{B\tilde{B}}|\psi^1\rangle_{RA\tilde{A}B\tilde{B}} = |\beta\rangle_{RB} \otimes |Q\rangle_{A\tilde{A}\tilde{B}}, \quad (2)$$

⁵ We always assume that these transformations act as the identities on the registers we do not specify explicitly.

where $|\beta\rangle_{RA} := (|00\rangle_{RA} + |11\rangle_{RA})/\sqrt{2}$ denotes an EPR pair on registers RA and $|P\rangle_{\tilde{A}B\tilde{B}}$ and $|Q\rangle_{A\tilde{A}\tilde{B}}$ are arbitrary pure states.

Lemma 5.1. *Let $|\psi^0\rangle, |\psi^1\rangle$ be states that fulfil (1) and (2). Then,*

$$|\langle\psi^0|\psi^1\rangle| \leq 1/2.$$

Proof. Multiplying both sides of (1) with $U_{A\tilde{A}}^\dagger$ and multiplying (2) with $V_{B\tilde{B}}^\dagger$, we can write

$$\begin{aligned} |\langle\psi^0|\psi^1\rangle| &= |\langle\beta|_{RA}\langle P|_{\tilde{A}B\tilde{B}} U_{A\tilde{A}} V_{B\tilde{B}}^\dagger |\beta\rangle_{RB}|Q\rangle_{A\tilde{A}\tilde{B}}| \\ &= |\langle\beta|_{RA}\langle P'|_{\tilde{A}B\tilde{B}}|\beta\rangle_{RB}|Q'\rangle_{A\tilde{A}\tilde{B}}| \\ &= |\langle P'|_{\tilde{A}B\tilde{B}}\langle\beta|_{RA}|\beta\rangle_{RB}|Q'\rangle_{A\tilde{A}\tilde{B}}|, \end{aligned}$$

where we used that $U_{A\tilde{A}}$ and $V_{B\tilde{B}}$ commute and defined $|P'\rangle_{\tilde{A}B\tilde{B}} := V_{B\tilde{B}}|P\rangle_{\tilde{A}B\tilde{B}}$ and $|Q'\rangle_{A\tilde{A}\tilde{B}} := U_{A\tilde{A}}|Q\rangle_{A\tilde{A}\tilde{B}}$. The last equality is just rearranging terms that act on different registers.

Note that writing out the partial inner product between $|\beta\rangle_{RA}$ and $|\beta\rangle_{RB}$ gives

$$\langle\beta|_{RA}|\beta\rangle_{RB} = \frac{1}{2}(\langle 0|_A|0\rangle_B + \langle 1|_A|1\rangle_B),$$

where the operator in the parenthesis “transfers” a qubit from register A to register B . Hence,

$$\begin{aligned} |\langle\psi^0|\psi^1\rangle| &= |\langle P'|_{\tilde{A}B\tilde{B}} \frac{1}{2}(\langle 0|_A|0\rangle_B + \langle 1|_A|1\rangle_B)|Q'\rangle_{A\tilde{A}\tilde{B}}| \\ &= \frac{1}{2} \cdot |\langle P'|_{\tilde{A}B\tilde{B}}|Q'\rangle_{B\tilde{A}\tilde{B}}| \\ &\leq \frac{1}{2}, \end{aligned}$$

where the last step follows from the fact that the inner product between any two unit vectors on the same registers can be at most 1. \square

5.2 Squeezing Many Vectors in a Small Space

For the sake of completeness, we reproduce here an argument similar to [NC00, Section 4.5.4] about covering the state space of dimension d with patches of radius ε .

Lemma 5.2. *Let \mathcal{B} be a set of 2^n distinct unit vectors in a complex Hilbert space of dimension d , with pairwise absolute inner product at most $1/2$. Then, the dimension d has to be in $\Omega(n)$.*

Proof. For any two vectors $|v\rangle, |w\rangle$, we can rotate the space such that $|v\rangle = |0\rangle$ and $|w\rangle = \cos\theta|0\rangle + \sin\theta|1\rangle$ for two orthogonal vectors $|0\rangle$ and $|1\rangle$. The *Euclidean distance* between $|v\rangle$ and $|w\rangle$ can be expressed as

$$\begin{aligned} ||v\rangle - |w\rangle| &= |(1 - \cos\theta)|0\rangle - \sin\theta|1\rangle| \\ &= \sqrt{(1 - \cos\theta)^2 + \sin^2\theta} \\ &= \sqrt{1 - 2\cos\theta + \cos^2\theta + \sin^2\theta} \\ &= \sqrt{2}\sqrt{1 - \cos\theta}. \end{aligned}$$

If $|v\rangle$ and $|w\rangle$ have absolute inner product at most $1/2$, we have that $|\cos\theta| \leq 1/2$ and hence $||v\rangle - |w\rangle| \geq 1$. Therefore, the vectors in \mathcal{B} have pairwise Euclidean distance at least 1. The set of unit vectors $|w\rangle$ with Euclidean distance at most δ from $|v\rangle$ is called *patch of radius δ around $|v\rangle$* . It follows that patches of radius $1/2$ around every vector in the set \mathcal{B} do not overlap.

The space of all d -dimensional state vectors can be regarded as the real unit $(2d-1)$ -sphere, because the vector has d complex amplitudes and hence $2d$ real degrees of freedom with the restriction that the sum of the squared amplitudes is equal to 1. Notice that the Euclidean distance between complex vectors $|v\rangle, |w\rangle$ remains unchanged if we regard these vectors as points of the real unit $(2d-1)$ -sphere.

The surface area of a patch of radius $1/2$ near any vector is lower bounded by the volume of a $(2d-2)$ -sphere of radius ε where ε is a constant slightly less than $1/2$.⁶ We use the formula $S_k(r) = 2\pi^{(k+1)/2}r^k / \Gamma((k+1)/2)$ for the surface area of a k -sphere of radius r , and $V_k(r) = 2\pi^{(k+1)/2}r^{k+1} / [(k+1)\Gamma((k+1)/2)]$ for the volume of a k -sphere of radius r . The total surface area of all patches, which is at least $2^n \cdot V_{2d-2}(\varepsilon)$, is not more than the total surface of the whole sphere $S_{2d-1}(1)$. Inserting the formulas, we get

$$2^n \cdot 2\pi^{d-\frac{1}{2}} \frac{\varepsilon^{2d-1}}{(2d-1)\Gamma(d-\frac{1}{2})} \leq 2\pi^d \frac{1}{\Gamma(d)}$$

Using the fact that $\frac{\Gamma(d-\frac{1}{2})}{\Gamma(d)} \leq \frac{1}{d}$, we conclude that

$$2^n \leq \sqrt{\pi} \left(2 - \frac{1}{d}\right) \varepsilon^{-(2d-1)} \leq 2\sqrt{\pi} \varepsilon^{-(2d-1)}.$$

As $\varepsilon < 1/2$, we obtain that d has to be in $\Omega(n)$. □

5.3 The Lower Bound

We consider perfect attacks on protocol PV_{qubit} from Figure 2. We allow the players one round of simultaneous quantum communication which we model as follows. Let $|\psi\rangle_{RA\tilde{A}A_C B\tilde{B}B_C}$ be the pure state after Alice received the EPR half from the verifier. The one-qubit register R holds the verifier's half of the EPR-pair, the one-qubit register A contains Alice's other half of the EPR-pair, the register \tilde{A} is Alice's part of the pre-shared entangled state and the register A_C holds the qubits that will be communicated to Bob. The registers $B\tilde{B}B_C$ belong to Bob where B holds one qubit and \tilde{B} is Bob's part of the entangled state and the B_C register will be sent to Alice. We denote by q_A the total number of qubits in registers \tilde{A} and A_C and by q_B the total number of qubits in \tilde{B} and B_C . The overall state is thus a unit vector in a complex Hilbert space of dimension $d := 2^{2+q_A+1+q_B}$.

In the first step of their attack, Alice performs a unitary transform U^x depending on her classical input x on her registers $A\tilde{A}A_C$. Similarly, Bob performs a unitary transform V^y depending on y on registers $B\tilde{B}B_C$. After the application of these transforms, the communication registers A_C and B_C and the classical inputs x and y are exchanged. A final unitary transform (performed either by Alice or Bob) depending on both x, y “unveils” the qubit either in Alice's register A or in Bob's register B .

Theorem 5.3. *Let f be injective for Bob. Assume that Alice and Bob perform a perfect attack on protocol PV_{qubit} . Then, the dimension d of the overall state (including the quantum communication) is in $\Omega(n)$.*

⁶ The patch is a “bended” version of this volume.

Proof. We assume that the player's actions are unitary transforms as described before the theorem.

We investigate the set \mathcal{B} of overall states after Bob performed his operation, but *before* Alice acts on the state. These states depend on Bob's input $y \in \{0, 1\}^n$,

$$\mathcal{B} := \left\{ V_{B\tilde{B}B_C}^y |\psi\rangle_{RA\tilde{A}A_CB\tilde{B}B_C} : y \in \{0, 1\}^n \right\}.$$

We claim that for any two different n -bit strings $y \neq y'$, the corresponding two vectors $V^y|\psi\rangle$ and $V^{y'}|\psi\rangle$ in \mathcal{B} have an absolute inner product of at most $1/2$.

Due to the injectivity of f , there exists an input x for Alice such that $f(x, y) \neq f(x, y')$. Applying Alice's unitary transform U^x to both vectors does not change their inner product, i.e.

$$|\langle \psi | (V^y)^\dagger V^{y'} | \psi \rangle| = |\langle \psi | (V^y)^\dagger (U^x)^\dagger U^x V^{y'} | \psi \rangle|.$$

As $f(x, y) \neq f(x, y')$, the qubit has to end up on different sides. Formally, there exist unitary transforms $K_{A\tilde{A}B_C}$ and $L_{B\tilde{B}A_C}$ that “unveil” the qubit in register A or B respectively. Hence, we can apply Lemma 5.1 to prove the claim that the two vectors $V^y|\psi\rangle$ and $V^{y'}|\psi\rangle$ have an absolute inner product of at most $1/2$. In particular, all of the vectors in \mathcal{B} are distinct. Applying Lemma 5.2 yields the theorem. \square

5.4 Functions For Which Perfect Attacks Need a Large Space

Theorem 5.4. *For any starting state $|\psi\rangle$ of dimension d , there exists a Boolean function on inputs $x, y \in \{0, 1\}^n$ such that any perfect attack on PV_{qubit} requires d to be exponential in n .*

Proof (sketch). We consider covering the sphere with K patches of vectors whose pairwise absolute inner product is larger than $\frac{\sqrt{3}}{2}$ (which corresponds to an Euclidean distance of $\varepsilon = \sqrt{2}\sqrt{1 + \sqrt{3}}/2 \approx 0.52$). This partitioning also induces a partitioning on all possible unitary operations of Alice and Bob. We say that two actions A and A' are in the same patch if they take the starting state $|\psi\rangle$ to the same patch. In other words, if two actions are in the same patch then

$$|\langle \psi | A'^\dagger A | \psi \rangle| \geq \frac{\sqrt{3}}{2}.$$

Claim. Given two actions of Alice A, A' coming from the same patch i , and two actions of Bob B, B' coming from the same patch j , the inner product between $BA|\psi\rangle$ and $B'A'|\psi\rangle$ has magnitude at least $\frac{1}{2}$.

Proof (of the claim). Since Alice and Bob act on different parts of the state, their actions commute. Write $|\psi_A\rangle := A'^\dagger A|\psi\rangle$ and $|\psi_B\rangle := B^\dagger B'|\psi\rangle$. Then the inner product can be written as

$$\langle \psi | A'^\dagger B'^\dagger B A | \psi \rangle = \langle \psi | B'^\dagger B A'^\dagger A | \psi \rangle = \langle \psi_B | \psi_A \rangle$$

Note that

$$|\langle \psi | \psi_A \rangle| = |\langle \psi | A'^\dagger A | \psi \rangle| \geq \frac{\sqrt{3}}{2},$$

so the angle θ between $|\psi_A\rangle$ and $|\psi\rangle$ is at most $\arccos \frac{\sqrt{3}}{2} = \frac{\pi}{6}$. The same holds for the angle between $|\psi_B\rangle$ and $|\psi\rangle$. We can upper bound the total angle between $|\psi_A\rangle$ and $|\psi_B\rangle$ by the sum of these angles, giving a total angle of at most $\frac{\pi}{3}$. This corresponds to a lower bound on the inner product of $\cos \frac{\pi}{3} = \frac{1}{2}$. \square

So there exists no pair of combined actions AB and $A'B'$, with A and A' in patch i and B and B' in patch j , such that the qubit ends up on Alice's side for AB and on Bob's side for $A'B'$. Therefore, the combination of i and j completely determines the destination of the qubit and hence the output of the function. If K denotes the number of patches, then there are K^{2^n} possible strategies for Alice and K^{2^n} possible strategies for Bob. Hence, the number of combined strategies (possibly resulting in different functions) is at most $K^{2 \cdot 2^n}$.

It is shown in [NC00, Section 4.5.4] that we need at least $K = \Omega(\frac{1}{\varepsilon^{d-1}})$ patches. Using the same counting argument as in Proposition 4.11, we have that

$$2^{2^{2^n}} \geq \Omega\left(\frac{1}{\varepsilon^{(d-1)2 \cdot 2^n}}\right),$$

from which follows that for some function, d has to be exponential in n . □

6 Conclusion and Open Questions

We defined the garden-hose model and gave first results for the analysis of a specific scheme for quantum position-based cryptography. This scheme only requires the honest prover to work with a single qubit, while the dishonest provers potentially have to manipulate a large quantum state, making it an appealing scheme to further examine. The garden-hose model captures the power of attacks that only use teleportation, giving upper bounds for the general scheme, and lower bounds when restricted to these attacks.

The garden-hose model is a new model of communication complexity, and there are many open questions in relation to this model. Can we find better upper and lower bounds for the garden-hose complexity of the studied functions? The constructions given in [Spe11] still leave a polynomial gap between lower and upper bounds for many functions. It would also be interesting to find an explicit function for which the garden-hose complexity is provably large, the counting argument in Proposition 4.11 only shows the existence of such functions.

Another relevant extension to our results is the examination of the randomized case: If we allow Alice and Bob to give the wrong answer with small probability, what are the lower and upper bounds we can prove in the garden-hose model? For example, assuming shared randomness between Alice and Bob, we can use results from communication complexity to show a large gap between the randomized garden-hose complexity of the equality function, and the deterministic garden-hose complexity of equality.

A possible interesting extra restriction on the garden-hose model would involve limiting the computational power of Alice and Bob. For example to polynomial time, or the output of quantum circuits of polynomial size. By bounding not only the amount of entanglement, but also the amount of computation with a realistic limit, perhaps stronger security proofs are possible.

On the other hand, one could also consider the quantum garden-hose model where Alice and Bob are additionally allowed to use shared entanglement. In order to keep the correspondence with quantum attacks on position-verification protocols, one would have to take into account both the number of required pipes and the size of the auxiliary entangled state the players use.

As a final question, we can ask: How does the protocol behave under parallel repetition? When executing the protocol once, the dishonest provers always have a large probability of cheating the verifiers; even the naïve method of measuring the qubit and distributing the result will work with a probability of at least 0.75. By using the protocol multiple times in parallel, given a situation where the adversaries have a small error, it might be possible to increase the probability that the dishonest provers are caught to arbitrarily close to 1. However, from complexity theory we know

similar situations where provers can achieve a lower error probability than expected on first sight. In our setting, it remains to be proven that we can always amplify the probability of the cheaters getting caught.

Acknowledgments

HB and FS are supported by an NWO Vici grant and the EU project QCS. CS is supported by an NWO Veni grant. We thank Louis Salvail for useful discussions about the protocol PV_{qubit} and Matthias Christandl for pointing out to us Lemma 3.2.

References

- Bar89. David A. Barrington. Bounded-width polynomial-size branching programs recognize exactly those languages in NC1. *Journal of Computer and System Sciences*, 164:150–164, 1989.
- BC94. Stefan Brands and David Chaum. Distance-bounding protocols. In *EUROCRYPT’93*, pages 344–359. Springer, 1994.
- BCF⁺11. Harry Buhrman, Nishanth Chandran, Serge Fehr, Ran Gelles, Vipul Goyal, Rafail Ostrovsky, and Christian Schaffner. Position-based quantum cryptography: Impossibility and constructions. In Phillip Rogaway, editor, *Advances in Cryptology CRYPTO 2011*, volume 6841 of *Lecture Notes in Computer Science*, pages 429–446. Springer Berlin / Heidelberg, 2011.
- BK11. Salman Beigi and Robert König. Simplified instantaneous non-local quantum computation with applications to position-based cryptography. *arXiv:1101.1065v1*, January 2011.
- Bus04. Laurent Bussard. *Trust Establishment Protocols for Communicating Devices*. PhD thesis, Eurecom-ENST, 2004.
- CCS06. Srdjan Capkun, Mario Cagalj, and Mani Srivastava. Secure localization with hidden and mobile base stations. In *IEEE INFOCOM*, 2006.
- CFG⁺10. Nishanth Chandran, Serge Fehr, Ran Gelles, Vipul Goyal, and Rafail Ostrovsky. Position-based quantum cryptography. *arXiv:1005.1750v2*, May 2010.
- CGMO09. Nishanth Chandran, Vipul Goyal, Ryan Moriarty, and Rafail Ostrovsky. Position based cryptography. In *CRYPTO 2009*, pages 391–407. Springer, 2009.
- CH05. Srdjan Capkun and Jean-Pierre Hubaux. Secure positioning of wireless devices with application to sensor networks. In *IEEE INFOCOM*, pages 1917–1928, 2005.
- IH08. Satoshi Ishizaka and Tohya Hiroshima. Asymptotic teleportation scheme as a universal programmable quantum processor. *Phys. Rev. Lett.*, 101(24):240501, Dec 2008.
- IH09. Satoshi Ishizaka and Tohya Hiroshima. Quantum teleportation scheme by selecting one of multiple output ports. *Phys. Rev. A*, 79(4):042306, Apr 2009.
- KMS11. Adrian Kent, William J. Munro, and Timothy P. Spiller. Quantum tagging: Authenticating location via quantum information and relativistic signaling constraints. *Phys. Rev. A*, 84:012326, Jul 2011.
- KMSB06. Adrian Kent, William Munro, Tomothy Spiller, and Raymond Beausoleil. Tagging systems, 2006. US patent nr 2006/0022832.
- LBZ02. Jay Lawrence, Časlav Brukner, and Anton Zeilinger. Mutually unbiased binary observable sets on N qubits. *Physical Review A*, 65(3):1–5, February 2002.
- LL11. Hoi-Kwan Lau and Hoi-Kwong Lo. Insecurity of position-based quantum-cryptography protocols against entanglement attacks. *Phys. Rev. A*, 83(1):012322, Jan 2011.
- LMT97. K.-J. Lange, Pierre McKenzie, and Alain Tapp. Reversible space equals deterministic space. In *Proceedings of Computational Complexity. Twelfth Annual IEEE Conference*, pages 45–50. IEEE Comput. Soc, April 1997.
- Mal10a. Robert A. Malaney. Location-dependent communications using quantum entanglement. *Phys. Rev. A*, 81(4):042319, Apr 2010.
- Mal10b. Robert A. Malaney. Quantum location verification in noisy channels, Apr 2010. *arXiv:1004.4689v1*.

- NC00. Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge university press, 2000.
- SP05. Dave Singelee and Bart Preneel. Location verification using secure distance bounding protocols. In *IEEE MASS'10*, 2005.
- Spe11. Florian Speelman. Position-based quantum cryptography and the garden-hose game. Master's thesis, University of Amsterdam, 2011.
- SSW03. Naveen Sastry, Umesh Shankar, and David Wagner. Secure verification of location claims. In *WiSe'03*, pages 1–10, 2003.
- VN04. Adnan Vora and Mikhail Nesterenko. Secure location verification using radio broadcast. In *OPODIS'04*, pages 369–383, 2004.
- ZLFW06. Yanchao Zhang, Wei Liu, Yuguang Fang, and Dapeng Wu. Secure localization and authentication in ultra-wideband sensor networks. *IEEE Journal on Selected Areas in Communications*, 24:829–835, 2006.